

ERGODIC INFINITE GROUP EXTENSIONS OF GEODESIC FLOWS ON TRANSLATION SURFACES

DAVID RALSTON AND SERGE TROUBETZKOY

ABSTRACT. We show that generic infinite group extensions of geodesic flows on square tiled translation surfaces are ergodic in almost every direction, subject to certain natural constraints. Recently K. Frączek and C. Ulcigrai have shown that certain concrete staircases, covers of square-tiled surfaces, are not ergodic in almost every direction. In contrast we show the almost sure ergodicity of other concrete staircases. An appendix provides a combinatorial approach for the study of square-tiled surfaces.

1. INTRODUCTION

In 1986 Kerchhoff, Masur and Smillie showed that the geodesic flow on a compact translation surface is ergodic in almost every direction [17]. On the other hand, the study of the dynamical properties on periodic infinite translation surfaces is in its infancy. There is a natural necessary and sufficient condition for the recurrence of the geodesic flow in almost every direction for \mathbb{Z} or \mathbb{R} covers [11]. Generic non-periodic translation surfaces are almost surely recurrent [29]. The almost sure diffusion rate in the periodic full occupancy wind-tree model, a certain important example of infinite translation surface, has also recently been understood [7]. We turn to the question of ergodicity. In [12], a natural example the “simple staircase” was shown to be ergodic in a.e. direction. In fact, the authors noticed that the ergodicity reduces to a classically studied question of ergodicity of cylinder flows (group valued skew products) over circle rotations initiated by Conze and studied by many authors, see [4] for a good survey. The question of ergodicity for other periodic translation surfaces reduces to the question of ergodicity of cylinder flows over interval exchange transformations.

Recently, Frączek and Ulcigrai have shown that certain periodic infinite translation surfaces are not ergodic in almost every direction [8]. More specifically, if the compact square-tiled translation surface M is in the stratum $\mathcal{H}(2)$, and the infinite translation surface \tilde{M} is a \mathbb{Z} unramified cover of M , then the infinite surface is not ergodic in a.e. direction. They showed a similar result for some \mathbb{Z}^2 extensions such as the periodic full occupancy wind-tree model.

In this article, we study the case when the surface \tilde{M} is a ramified cover of M . This is generically the case. In the language of ergodic theory the flow on \tilde{M} is a G -valued skew product over the geodesic flow on M for some Abelian locally compact group G . Our result is a sufficient condition for both recurrence and almost sure

Date: January 19, 2012.

The first author is supported by the Center for Advanced Studies at Ben Gurion University of the Negev as well as the Israel Council for Higher Education, and was partially supported by the Erwin Schrödinger International Institute for Mathematical Physics during preparation of this manuscript. This research is partially supported by the ANR project Perturbations.

ergodicity to hold. This condition is based on ideas developed by M. Boshernitzan [2] and K. Schmidt [26].

While our theorem is somewhat technical to state, it is applicable to many concrete examples. Suppose that M is a square tiled surface with a single-cylinder direction (for example the torus), then any reasonable method of producing a \mathbb{Z}^d (or more generally an Abelian locally compact group G) infinite surface \tilde{M} which is a skew product over M is generically ergodic in almost every direction. We also produce concrete examples of staircases ergodic in almost every direction such as the staircase in Figure 1. In the appendix, we give a procedure for determining which square-tiled surfaces have a single-cylinder direction, as well as a quick method for constructing such surfaces and determining their genus as well as the number and type of singularities.




FIGURE 1. The periodic staircase with alternating steps of length 2 and 3 is ergodic in almost every direction. Opposite sides are identified.

J. Chaika and P. Hubert have recently announced the following related result [3]: for a given $f : [0, 1) \rightarrow \mathbb{R}$ with integral 0, which is the finite sum of characteristic functions of intervals, for a set of full measure of interval exchange transformations T the skew product $T_f(x, i) = (Tx, i + f(x))$ into the closed subgroup of \mathbb{R} generated by the values of $f(x)$ is ergodic. Their methods of proof are close to ours.

2. TRANSLATION SURFACES

In this section we will give background material on translation surfaces as well as some definitions for later use, see [15, 21, 31] for details. A translation surface is a 2-dimensional manifold M and a finite set of points $D = \{d_1, d_2, \dots, d_m\}$ and an open cover of $M \setminus D$ by sets $\{U_\alpha\}$ together with charts $\phi_\alpha : U_\alpha \rightarrow \mathbb{R}^2$ such that for all α, β , with $U_\alpha \cap U_\beta \neq \emptyset$, $\phi_\alpha \circ \phi_\beta^{-1}$ is a translation on its domain of definition and at each singular point $d_i \in D$ the surface has a $2\pi n$ cone singularity for some integer $n \geq 1$.

The surface M has a flat metric obtained by pulling back the Euclidean metric on the plane via the coordinate charts. This metric is defined away from the set D . In this metric geodesics that do not go through singularities project via the

charts to straight lines in the plane in a fixed direction, and such geodesics are either periodic or simple (do not intersect themselves). If the geodesic in direction θ starting at $x \in M$ avoids D then we define the geodesic flow $\varphi_t^\theta(x)$ to the point obtained after moving in the direction θ for a time t starting at x . Since θ is usually fixed, we will often suppress the θ dependence of the flow and simply write $\varphi_t(x)$. If the geodesic hits D , then the geodesic flow $\varphi_t(x)$ is defined up to this time. The geodesic flow preserves the surface area measure μ .

Let $\Gamma \subset SL_2(\mathbb{R})$ denote the *Veech group* of the surface, the group of affine transformations (seen in $SL_2(\mathbb{R})$ as acting on the charts of M) which preserve the structure of M , and therefore permute the singularities D . If $SL_2(\mathbb{R})/\Gamma$ is of finite volume (i.e. Γ is a *lattice*), we say that M is a *Veech surface*.

Throughout the article we will assume that M is a compact translation surface; by [17] the geodesic flow is (uniquely) ergodic in almost every θ . We will form *skew products* over these flows in the following manner: let G be a locally compact Abelian group with translation-invariant metric $\|\cdot\|$, whose identity element we denote 0, with Haar measure ν , and for a finite collection of geodesic path segments $\{\gamma_i\}$ (disjoint without loss of generality) whose union we denote γ , let $f : \gamma \rightarrow G$ be constant on each γ_i and 0 elsewhere. Given a point $x \in M$ and a direction θ which is not parallel to any γ_i , denote

$$S_t(x) = \sum_{\substack{0 \leq s < t \\ f(\varphi_s(x)) \neq 0}} f(\varphi_s(x)),$$

$$S_t^{-1}(g) = \{x \in M : S_t(x) = g\}, \quad B_\epsilon(g) = \{y \in G : \|y - g\| < \epsilon\}.$$

Then the transformation $\tilde{\varphi}_t$, from $\tilde{M} = M \times G$ to itself, is given by

$$\tilde{\varphi}_t(x, g) = (\varphi_t(x), g + S_t(x)),$$

and it is this flow which we study; $\tilde{\varphi}_t$ preserves the measure $\tilde{\mu} = \mu \times \nu$. An *essential value* of a general skew product $\{\tilde{M}, \tilde{\mu}, \tilde{\varphi}_t\}$ is some $g \in G$ such that for every $\epsilon > 0$ and $A \subset M$ with $\mu(A) > 0$, there is some $t > 0$ such that

$$\mu(A \cap \varphi_t(A) \cap S_t^{-1}(B_\epsilon(g))) > 0.$$

We denote by $E(\tilde{\varphi}_t)$ the union of all essential values. Then if $E(\tilde{\varphi}_t)$ is nonempty, it is a closed subgroup of G , and $\{\tilde{M}, \tilde{\varphi}_t\}$ is ergodic if and only if $E(\tilde{\varphi}_t) = G$ and $\{M, \varphi_t\}$ is ergodic (see [1] for example). If $0 \in E(\tilde{\varphi}_t)$, then the system is called *recurrent*.

Definition 2.1. We will call a sequence of sets C_n a *quasi-rigidity sequence of sets* in the direction θ if there is a sequence $t_n \rightarrow \infty$, a sequence $\epsilon_n \rightarrow 0$ and a fixed $\epsilon > 0$ such that

- $\mu(C_n) \geq \epsilon$,
- $\forall x \in C_n, \|\varphi_{t_n}(x) - x\| \leq \epsilon_n$, and
- $\mu(C_n \triangle \varphi_1 C_n) \leq \epsilon_n$.

That is, each set is of nontrivial measure, the map φ_{t_n} acts almost like the identity on these sets, and they are nearly invariant under the flow φ_1 (the use of φ_1 is arbitrary; any fixed positive time would do).

Lemma 2.2. *Suppose that $\{M, \varphi_t\}$ is ergodic and that for some $g \in G$ we have a quasi-rigidity sequence of sets C_n such that*

$$\lim_{n \rightarrow \infty} \sup_{x \in C_n} \|S_{t_n}(x) - g\| = 0.$$

Then $g \in E(\tilde{\varphi}_t)$.

Proof. This is merely a restatement for continuous flows of [5, Cor. 2.6]. In that reference, it is assumed that $G = \mathbb{R}^d$, but the techniques apply to any locally compact Abelian group with Haar measure and translation-invariant metric, as we have assumed. A similar remark appears as [16, Thm. 11]; the notion of restricting a search for essential values to certain subsets whose measure is bounded away from zero is a common step in the study of periodic group extensions of finite-measure systems. \square




FIGURE 2. Decomposition of $\{M, \varphi_t\}$ into periodic cylinders in direction θ_i of widths $\alpha_{i,j}$ and heights $h_{i,j}$; the top and bottom of each rectangle are identified, and φ_t flows vertically at unit speed. The error compared to the flow in direction θ relative to the width of the cylinders is bounded for periodic approximations, bounded away from zero for fairly good periodic approximations, and converging to zero for good periodic approximations.

Definition 2.3. If the geodesic flow φ_t in a particular direction θ decomposes the space M into a finite union of sets on which φ_t is periodic (see Figure 2), the direction θ is said to *admit a representation by periodic cylinders*. Given a sequence $\{\theta_i\}$ of directions which admit a representation by periodic cylinders, denote the cylinders by $A_{i,j}$, the width of $A_{i,j}$ by $\alpha_{i,j}$, and the height of each cylinder (i.e. the period of the geodesic flow in this cylinder) by $h_{i,j}$.

The quantity

$$E_i = \max \left\{ \frac{h_{i,j} \tan |\theta - \theta_i|}{\alpha_{i,j}} : j = 1, 2, \dots, n_i \right\}$$

will be called the *error of the approximation θ_i* . If $\theta_i \rightarrow \theta$ and $E_i \in O(1)$, we say that $\{\theta_i\}$ is a *sequence of periodic approximations* to the flow in the direction θ . Note that if $\theta_i \rightarrow \theta$ and θ does *not* admit a representation by periodic cylinders (and only countably many directions, those of *saddle connections*, i.e. orbit segments which

begin and end at singular points, may have such a representation), then necessarily we have

$$\lim_{i \rightarrow \infty} \min_j h_{i,j} = \infty.$$

If furthermore $E_i \rightarrow 0$ we call the sequence $\{\theta_i\}$ a *good sequence of periodic approximations*. If there is some $\delta > 0$ such that for all i we have $E_i \in (\delta, 1/2)$ we will call the sequence a *fairly good sequence of periodic approximations*. See again Figure 2. If $\inf \{\mu(A_{i,j}) = h_{i,j} \cdot \alpha_{i,j} : i = 1, 2, \dots, j = 1, 2, \dots, n_i\} > 0$, we will say that this sequence *has cylinders of comparable measures*.

Lemma 2.4. *Suppose that θ admits a sequence of good periodic approximations $\{\theta_i\}$ with cylinders of comparable measures, and let $\{i_k\}$ be an increasing sequence of positive integers. For some fixed $\delta > 0$, consider a sequence of subcylinders C_k ; connected subsets of the cylinder $A_{i_k,j}$ (for some j) which are invariant under the flow in direction θ_{i_k} , and whose measure is at least $\delta \mu(A_{i_k,j})$. Assume further that the boundary of each C_k is separated from the boundary of $A_{i_k,j}$ by a distance of at least $\delta \alpha_{i_k,j}$. Then the $\{C_k\}$ form a quasi-rigidity sequence of sets.*

Proof. That $\mu(C_k)$ is bounded away from zero follows from the condition that the cylinders have comparable measures. As the sequence θ_i was taken to be a sequence of good periodic approximations, the relative error converges to zero. It now follows from the fact that each C_k is bounded away from the boundary of $A_{i_k,j}$ that for each $x \in C_k$, the orbit through time $h_{i_k,j}$ of x remains in $A_{i_k,j}$, from which it follows that

$$\|\varphi_{h_{i_k,j}}(x) - x\| \leq E_{i_k} \alpha_{i_k,j} \xrightarrow{k \rightarrow \infty} 0.$$

Finally, we have

$$\mu(\varphi_1 C_k \triangle C_k) \leq 2 \cdot E_{i_k} \alpha_{i_k,j}$$

(the difference is exactly those x so close to the boundary of C_k that they flow out of the subcylinder in unit time), which certainly converges to zero. \square

Definition 2.5. Suppose that the geodesic flow in direction θ admits a representation by a single periodic cylinder. Then θ will be called a *single-cylinder direction*.

Single-cylinder directions are of great convenience in studying translation surfaces. It was shown in [13] that if M is square-tiled with a prime number of tiles and in $\mathcal{H}(2)$ (genus two with a single singularity), then M has a single-cylinder direction. In [22, Cor. A.2] McMullen extended this result to all square-tiled surfaces in $\mathcal{H}(2)$. Kontsevich and Zorich showed that there is a dense set of square tiled surfaces in any stratum with single cylinder directions [18, Lemma 18]. Finally Laneeau and Nguyen have identified certain substrata of $\mathcal{H}(4)$ and $\mathcal{H}(6)$ which have single cylinder directions [19]. See Appendix A for a general technique for constructing square-tiled surfaces with a single-cylinder direction.

The following is a standard argument in the study of translation surfaces, but hardly elementary:

Lemma 2.6. *Suppose that $\{M, \varphi_t\}$ has at least one direction θ' which admits a representation by N periodic cylinders, and suppose that M is a Veech surface with Veech group Γ . Then for almost every direction θ , there is a sequence $\{\theta_i\}$ of directions, chosen from the directions $\Gamma(\theta')$, which are a good sequence of periodic approximations, with N cylinders of comparable measures, to the flow in the direction θ . Furthermore, for almost every direction θ there is a sequence $\{\theta_i\}$ taken*

from the directions $\Gamma(\theta')$ of fairly good periodic approximations with N cylinders of comparable measures.

Proof. Any element $\gamma \in \Gamma$ will transform the flow in direction θ' to a flow in another direction which must necessarily also admit a representation by periodic cylinders. The number of cylinders and their measures are not changed by the action of $\gamma \in SL(2, \mathbb{R})$, so all that must be shown is that for almost every θ , a sequence may be chosen out of this orbit which are a *good* or *fairly good* sequence of periodic approximations to the flow in the direction of θ .

The assumption that M is a Veech surface ensures that Γ , the Veech group, is *Fuchsian of the first kind*. The fact that almost-every direction has a sequence of good approximations coming from the orbit under Γ of the single-cylinder direction follows directly from [23, §9] (see [28] for a more general case). We state the application of this theorem to our setting in the language of this article: for every $\varepsilon > 0$, there exists a set of full measure $\Theta(\varepsilon) \subset \mathbb{S}^1$ such that for any $\theta \in \Theta(\varepsilon)$, there exists a sequence of N cylinder directions with slopes p_n/q_n such that $|\theta - p_n/q_n| \leq \varepsilon/q_n^2$. Similarly as in [14], the conclusions for good approximations follow from Patterson's theorem. The appropriate lower bound for fairly good approximations follows from the ergodicity techniques used to prove similar statements in [16, §2.2]: the orbit of a completely periodic direction under Γ will intersect arbitrary sets of positive measure infinitely often in $SL_2(\mathbb{R})/\Gamma$. This quotient is of finite volume and has finitely many cusps by assumption, so a nontrivial annular neighborhood around a generic slope will also contain infinitely many points from the orbit $\Gamma(\theta')$; the use of a fixed annulus as opposed to a sequence of shrinking neighborhoods distinguishes fairly good approximations from good approximations. \square

In particular, if M is a Veech surface, then if the flow in direction θ' is not periodic the surface admits a representation by periodic cylinders of comparable measures with whatever strength in approximation we desire. The conclusion of Lemma 2.6 can in fact be strengthened to state that the set of exceptional directions (those which cannot be approximated to the desired accuracy from the orbit under Γ of any given single-cylinder direction) is not just of measure zero, but indeed of Hausdorff dimension less than one. See again [16] for similar statements; we state only that almost-every direction can be so approximated in the interest of simplicity and generic appeal.

3. GENERIC CONSTRUCTIONS

Definition 3.1. Let θ , the direction of the geodesic flow, be fixed such that $\{M, \varphi_t\}$ admits a sequence of periodic approximations $\{\theta_i\}$. Let \tilde{D} denote an arbitrary finite subset of M . Then a point $x \in (\tilde{D} \setminus D)$ is said to *avoid* $D \cup \tilde{D}$ in direction θ if for some fixed $\epsilon > 0$ there is a sequence $\{\epsilon_i\}$ so that each of the following holds:

- for each i , x belongs to the interior of some $A_{i,j}$,
- the ϵ_i -neighborhood around the $h_{i,j}$ sec $|\theta_i - \theta|$ orbit of x in the direction of θ remains in $A_{i,j}$ (restricted to those j for which $x \in A_{i,j}$) and does not intersect $D \cup (\tilde{D} \setminus x)$, and
- $\inf_{i \rightarrow \infty} \epsilon_i h_{i,j} \geq \epsilon$.

A sequence of neighborhoods C_i for which x satisfies this definition for the same $\epsilon > 0$ will be called a *sequence of orbit-neighborhoods* for x . See Figure 4 for an illustration (the sets $C_{i,1}$ and $C_{i,2}$ will be defined later). If every point in \tilde{D} avoids

$D \cup \tilde{D}$, and if $D \cap \tilde{D} = \emptyset$, then we will say that \tilde{D} is a *self-avoiding set in direction* θ .

Note that we merely need *some* sequence of periodic approximations for each θ ; in situations where a canonical method for finding approximations exists, a subsequence of this canonical set suffices.

Proposition 3.2. *Suppose θ has a good sequence of periodic approximations and fix $P \in \mathbb{N}$. Then almost every element $\tilde{D} \in M^P$ (with respect to the product measure) is a self-avoiding set in the direction θ .*

Proof. It suffices to show that the proposition is true for $P = 1$; larger sets may then be inductively constructed one point at a time by presuming x_1 through x_{P-1} are in D for the purposes of finding x_P . Let θ be fixed such that $\{\theta_i\}$ is a good sequence of periodic approximations. For a fixed $\epsilon > 0$, let $B_i = B_i(\epsilon)$ be the complement of the ϵ' -neighborhood of the boundaries of the cylinders, where

$$\epsilon' = \frac{\epsilon \min_j \{\alpha_{i,j}\}}{2},$$

so that $\mu(B_i) \geq (1 - \epsilon)$. See Figure 3. Now note that if x belongs to infinitely many $B_i(\epsilon)$, it is self-avoiding singleton: for sufficiently large i we have the error $E_i < \epsilon$. As

$$\mu \left(\limsup_{i \rightarrow \infty} B_i \right) \geq \limsup_{i \rightarrow \infty} \mu(B_i),$$

for this choice of ϵ the set of self-avoiding singletons in direction θ is of measure at least $(1 - \epsilon)$, and ϵ was arbitrary. \square

Corollary 3.3. *Suppose that almost every θ admits a sequence of good approximations. Then almost every subset $\tilde{D} \in M^P$ is self-avoiding in almost every direction.*

Proof. We certainly have almost surely that $\tilde{D} \cap D = \emptyset$. For a full-measure set of θ , a full-measure set of \tilde{D} is self-avoiding in direction θ . Apply Fubini's theorem to reverse the order of quantifiers. \square




FIGURE 3. The ϵ -proportion removed in constructing $B_i(\epsilon)$.

Let $S = \{s_1, \dots, s_n\}$ be a finite subset of M (possibly including points in D), $L_i > 0$ and $\tau_i \in \mathbb{S}^1$ ($i = 1, \dots, n$) be for the moment arbitrary. Let γ_r be a geodesic of length $L_r > 0$ in direction τ_r originating at s_r . Note that we write “a” geodesic since if $s_r \in D$ there may be several possible such geodesics and if a geodesic originating at s_r arrives at a point of D , there are multiple possible continuations. Although generically this will not happen, a priori we allow all such geodesic segments. Let $a_{r,1}, \dots, a_{r,N(r)}$ be a finite set of points in γ_r ($N(r)$ is arbitrary but finite); assume that $a_{r,N(r)} = \varphi_{L_r}(s_r)$ (the terminus of γ_r) and set $a_{r,0} = s_r$. Denote by A the (finite) collection of all $a_{r,j}$ together with any $d \in D$ such that $d \in \gamma_i$ for some i .

Corollary 3.4. *For fixed n , for almost every construction of γ_r and A (with initial points s_i chosen according to a probability measure on M absolutely continuous with respect to μ , lengths l_r chosen according to some probability measure absolutely continuous with respect to Lebesgue on \mathbb{R}^+ and directions τ_r by a measure absolutely continuous with respect to Lebesgue measure on \mathbb{S}^1 , and A then chosen according to a measure on γ which is absolutely continuous with respect to Lebesgue measure), the set A is a self-avoiding set.*

Proof. We apply Proposition 3.2, noting simply that we certainly have almost-surely that no γ_i crosses a singularity $d \in D$. \square

Let G be an Abelian locally compact group with translation invariant metric endowed with Haar measure ν , and we will use the segments γ_i to define our skew product $\{\tilde{M}, \tilde{\varphi}_t\}$ as in §2. Define a function $f : M \rightarrow G$ to be piecewise constant on each γ_i with discontinuities a subset of A ; denote the value of f between $a_{i,j-1}$ and $a_{i,j}$ by $f_{i,j}$. Finally, denote

$$\sigma_{i,j} = \begin{cases} f_{i,1} & (j = 0) \\ f_{i,j+1} - f_{i,j} & (j = 1, 2, \dots, N(i) - 1) \\ -f_{i,N(i)} & (j = N(i)). \end{cases}$$

The function f is defined to be identity off the segments γ_i ; the values $\sigma_{i,j}$ are the ‘jumps’ seen in f while flowing along the path γ_i .

Definition 3.5. Let θ be fixed. A point $x \in (\tilde{D} \setminus D)$ which avoids $(\tilde{D} \cup D)$ with orbit neighborhoods C_k of heights $h_k \rightarrow \infty$ will be called *essential in direction θ* if there is some P such that along an infinite subsequence of orbit neighborhoods we have

$$\sup_{k=1,2,\dots} \sup_{y \in C_k} |S_{h_k}(y)| \leq P.$$

Showing that points are essential is in general a very demanding requirement. Single-cylinder directions, however, provide a ready approach.

Proposition 3.6 (Koksma-type inequality). *Suppose that M is a translation surface and the function f is defined as above. Assume further that θ' is a single cylinder direction, with the cylinder height given by h (and width $1/h$). Suppose that the flow in the single cylinder direction θ' , sees an average value of the identity in G . Then, for every x not on the boundary of the cylinder we have*

$$\|S_h(x)\| \leq 4 \sum_{j=1}^{N(i)} \sum_{i=1}^n \|f_{i,j}\|,$$

where the sum $S_h(x)$ is with respect to the flow in the direction θ' .

Proof. We denote by $q_{i,j}$ the projected length of the geodesic segment connecting $a_{i,j-1}$ to $a_{i,j}$ in the direction transverse to θ' . That the average value of the sums in the cylinder is identity is equivalent to

$$(1) \quad \sum_{j=1}^{N(i)} \sum_{i=1}^n f_{i,j} q_{i,j} = 0.$$

This condition is sufficient for recurrence of the flow in the event that $G = \mathbb{Z}$ or \mathbb{R} (see [25]), and is clearly necessary in all cases where G contains elements of infinite order.

Given a particular segment, the points $a_{i,j-1}$ and $a_{i,j}$ are in the cylinder (possibly on the boundary), and the segment between them completely crosses the cylinder some number of times. Denote this number by $P_{i,j}$. There are also possibly two partial crossings corresponding to the segment leaving $a_{i,j-1}$ and arriving into $a_{i,j}$ (in $P_{i,j} = 0$ then there is only one such segment); denote by $p'_{i,j}$ and $p''_{i,j}$ respectively the lengths of these two segments when measured in the transverse direction to the flow. We have

$$0 \leq p'_{i,j}, p''_{i,j} < \frac{1}{h}, \quad 0 \leq q_{i,j} - \frac{P_{i,j}}{h} \leq \frac{2}{h},$$

the latter of which we rewrite as

$$|P_{i,j} - q_{i,j}h| \leq 2.$$

As the entire space is shown in this single periodic cylinder, the entire segment is represented by these pieces, so (1) translates to

$$(2) \quad \sum_{j=1}^{N(i)} \sum_{i=1}^n f_{i,j} \left(\frac{P_{i,j}}{h} + p'_{i,j} + p''_{i,j} \right) = 0.$$

The periodic orbit of any point x in the direction θ' (ignoring the orbits containing the points $a_{i,j}$) crosses the segment from $a_{i,j-1}$ to $a_{i,j}$ a total of $P_{i,j} + \phi_{i,j}(x)$ times, where

$$\phi_{i,j}(x) \in \{0, 1, 2\},$$

according to whether it crosses neither ‘partial segment,’ one of them, or both. Then we have

$$\begin{aligned} S_h(x) &= \sum_{j=1}^{N(i)} \sum_{i=1}^n f_{i,j} (P_{i,j} + \phi_{i,j}(x)) \\ \|S_h(x)\| &\leq \left\| \sum_{j=1}^{N(i)} \sum_{i=1}^n f_{i,j} P_{i,j} \right\| + 2 \sum_{j=1}^{N(i)} \sum_{i=1}^n \|f_{i,j}\| \\ &\leq h \left\| \sum_{j=1}^{N(i)} \sum_{i=1}^n f_{i,j} q_{i,j} \right\| + 4 \sum_{j=1}^{N(i)} \sum_{i=1}^n \|f_{i,j}\| \\ &= 4 \sum_{j=1}^{N(i)} \sum_{i=1}^n \|f_{i,j}\|. \end{aligned} \quad \square$$

Theorem 3.7. *Let $a = a_{j,i}$ be an essential point with respect to a good sequence of periodic approximations. Then $\sigma = \sigma_{i,j}$ is an essential value of the skew product $\{\tilde{M}, \tilde{\varphi}_t\}$ for the flow in direction θ . In particular, the skew product is recurrent.*

Proof. Using a good sequence of periodic approximations allows us to consider for almost every x , for sufficiently large i the flow in direction θ and the flow in direction θ_i are indistinguishable through length $h_{i,j}$, at least insofar as the sums $S_{h_{i,j}}(x)$ are concerned. The point a has a sequence of orbit-neighborhoods C_k which remain in the same periodic cylinder $A_{i_k,j}$. The only discontinuity of f in C_k is the point a by definition. So we split the orbit-neighborhood into two sets of equal measure $C_{k,1}$ and $C_{k,2}$ so that the sum $S_{h_{i_k}}(x)$ is constant on each. Furthermore, the difference between the two values taken is exactly σ . See Figure 4. Without loss of generality, then, assume that for $x \in C_{k,1}$ and $y \in C_{k,2}$ we have

$$S_{h_{i_k,j}}(y) = S_{h_{i_k,j}}(x) + \sigma.$$

As there is a uniform bound P across all values and G is locally compact, by passing to a subsequence (which we continue to denote $\{i_k\}$) we may find some limit

$$\lim_{k \rightarrow \infty} S_{h_{i_k}}(y \in C_{k,1}) = g, \quad \lim_{k \rightarrow \infty} S_{h_{i_k}}(y \in C_{k,2}) = g + \sigma.$$

Both $C_{k,1}$ and $C_{k,2}$ are a quasi-rigidity sequence sets by Lemma 2.4, so by Lemma 2.2, both g and $g + \sigma$ are in $E(\tilde{\varphi}_t)$. As the set of essential values is a group, $\sigma \in E(\tilde{\varphi}_t)$. \square




FIGURE 4. The two halves of the orbit-neighborhood $C_i = C_{i,1} \cup C_{i,2}$ which do not intersect any discontinuities of f and whose sums differ by exactly σ in the proof of Theorem 3.7.

Corollary 3.8. *If the closed subgroup of G generated by the ‘jumps’ $\sigma_{j,i}$ at essential points $a_{j,i}$ is all of G , then the geodesic flow on \tilde{M} is ergodic in direction θ provided the underlying flow $\{M, \varphi_t\}$ is ergodic in direction θ . If all discontinuities $a_{i,j}$ are essential points, then every value $f_{i,j}$ is an essential value for $\{\tilde{M}, \tilde{\varphi}_t\}$.*

That is, for a generic method of making ‘cuts’ γ_i on the surface M and using them to form a G -fold cover, as long as the values $f_{i,j}$ include generators for the

group G , the extension is ergodic in almost every direction *provided we may show the discontinuities are essential points in almost-every direction*.

Example 3.9. Let M be the two-torus $\mathbb{T} = [0, 1)^2$ with standard identifications (which is flat without any singularities). Almost every direction θ has a good sequence of single-cylinder periodic directions given by a subsequence of the continued fraction convergents to θ , and the flow $\{M, \varphi_t\}$ is ergodic in any irrational direction. Let $\{g_i\}$ generate the group G , and for each i construct a pair of parallel geodesics of equal length, at random starting points, in random directions, and of random lengths (with the relevant distributions are all absolutely continuous with respect to Lebesgue). Define f to be $\pm g_i$ on these pairs. Then in any direction not parallel to such a segment, the projection of f to the base segment $[0, 1) \times \{0\}$ will be of bounded variation ($\text{Var}(f) = 4 \sum |g_i|$) and have integral zero, so by a combination of Propositions 3.2 and 3.6, Theorem 3.7, and Corollary 3.8, for generic such constructions the skew product $\{\tilde{M}, \tilde{\varphi}_t\}$ is ergodic in almost every direction.

Instead of generating ‘parallel pairs’ of cuts, we could generate *single* cuts which are subdivided appropriately. For example, suppose that $a_i g_i + b_i g'_i = 0$, where a_i and b_i are in \mathbb{R} , and the smallest closed subgroup in G containing all of the $\{g_i, g'_i\}$ is G itself. Generate i cuts of random length in random directions, and subdivide each cut into two portions whose ratio is a_i/b_i ; let f take value g_i on the segment of proportional length b_i and value g'_i on the other.

Note that any ‘reasonable’ method for generating cuts on M which ensures that the transfer function f has zero integral in almost-every direction will also generically ensure that the skew product $\{\tilde{M}, \tilde{\varphi}_t\}$ is ergodic in almost every direction, provided that bounds on sums in cylinders may be established; by Proposition 3.6 Veech surfaces which admit a single-cylinder direction will automatically obey such a bound.

Example 3.10. Let M be *square-tiled*: a finite cover of the torus such that M may be considered to be tiled by squares which share corners. Then M is a Veech surface; this fact was first shown by Veech using Boshernitzan’s criterion in [30], but modern proofs center around the use of arithmetic subgroups of $\text{SL}_2(\mathbb{R})$ (see e.g. [13, App. C], [9, Thm. 5.5]). Assume further that M has a single-cylinder direction, so that almost every direction θ admits a good sequence of periodic single-cylinder directions. By the same argument as in the previous example, then, for a generic method of constructing f the corresponding flow is ergodic in almost every direction.

See Appendix A for material relating to determining which square-tiled surfaces have a single-cylinder direction as well as for a simple method of generating such surfaces. While convenient, the assumption that M admits a single-cylinder direction is not necessary to apply the ideas and results of this section; an example of a class of surfaces with no single-cylinder directions for which certain constructions generically produce ergodic covers is addressed in Example 4.6 and the remark following that example.

We remark that the idea of comparing ergodic sums on two different sides of a discontinuity goes back at least as far as [26] and can be found more recently in [6], and the notion of finding generic points bounded a positive proportional distance from the orbits of singularities can be found in [2]. A similar approach to

studying ergodicity of real-valued skew products over rank-one interval exchange transformations can be found in [3].

4. GENERALIZED STAIRCASES

In this section we shall apply the same principles as in our generic constructions to certain *specific* skew products of square-tiled surfaces called *generalized staircases*.

Definition 4.1. Let M be a translation surface, and let \tilde{M} be a \mathbb{Z}^d cover of M formed by taking for each $i = 1, 2, \dots, d$ a pair of parallel cuts of equal length, with $f = \pm e_i$ on each cut, where e_i is the i -th standard basis element for \mathbb{Z}^d . If all of the cuts intersect at most at their endpoints, then \tilde{M} will be called a *generalized staircase* over M ; there is a natural visualization of the flow on the surface with some number of cuts which transfer the flow along a \mathbb{Z}^d -periodic cover. If furthermore M is square-tiled and each cut is of length one (the length of the sides of the squares) and all endpoints are integer points, \tilde{M} will be called a *natural staircase* over M .

Note that if each cut is nontrivial in homology, then the \mathbb{Z}^d -cover of M is *unramified*, and if $M \in \mathcal{H}(2)$, by [8, Thm. 1.4] the system $\{\tilde{M}, \tilde{\varphi}_t\}$ is *not* ergodic in almost every direction. However, if some of the cuts have endpoints which are *not* in D (the singularities of the flat metric on M), then we have constructed a ramified cover over a set of marked points D . It is direct to see that if M consists of a single square then only trivial natural staircases are possible. If we represent the torus with *two* squares, however, then a natural staircase which is a ramified cover becomes possible. This cover of the torus was presented in [12], where it was shown to be equivalent to a previously studied *cylinder transformation*, known to be ergodic in all irrational directions.

Generalized staircases can be parameterized by the base points and lengths of their respective cuts. The results of the previous section may be readily adapted under the addition restriction that the cuts are parallel to sides to ensure that generic generalized staircases are ergodic in almost every direction. However specific directions can have quite different behavior.

Example 4.2. Let $M = \mathbb{T}$ be the two-torus. Suppose that θ is a direction which admits a good sequence of periodic approximations. Then there exist uncountably many choices of β, r such that with cuts as in Figure 5, the skew product model of the corresponding generalized staircase $\{\tilde{\mathbb{T}}, \tilde{\varphi}_t\}$ is *nonregular* (the only essential values are $\{0, \infty\}$). This staircase is the same as the skew product considered by Conze in [4, Thm. 4.2] where exactly this conclusion is shown. Therefore the geodesic flow is *recurrent* (0 is an essential value), but not ergodic (not all of \mathbb{Z} are essential values).

The situation of natural staircases is in some sense more interesting: the severe restrictions imposed by the integer endpoints of the cuts present an immediate obstacle to directly applying the techniques of the previous section without modification. If the staircase is given by an unramified cover, then all endpoints of the cuts are singularities, and there is no adapting our techniques (indeed, for $M \in \mathcal{H}(2)$ the flow will not be ergodic in almost every direction by [8]). So assume that at least one of the cuts has an endpoint (with integer coordinates, as we consider a natural staircase) which is not a singularity. In a square-tiled surface, the directions




FIGURE 5. A generalized staircase over the 2-torus which is non-regular for uncountably many parameters β, r .

which admit representation by periodic cylinders correspond to rational-slope flows on the plane. The endpoints of our cuts are all integer points in the plane, as are all of the singularities \tilde{D} ; the endpoints of our cuts will always belong to the edge of the periodic cylinders and therefore orbit into singularities in the directions of the periodic approximations, a rather insurmountable impediment to directly apply Definition 3.1.

Thus far, however, the only identifications of the edges of the cylinders we have used is the identification of the top to the bottom, but there are also certain identifications on the *sides* of the cylinders. It is illustrative to make explicit which directions for square-tiled surfaces form different types of periodic approximation. Let M be square-tiled on k squares. The flow in direction θ induces a rotation by $\theta' = 1/\theta \pmod{1}$ on \mathbb{S}^1 when the bases of all squares are projected to a single circle. Let the continued fraction expansion of θ' be given by

$$\theta' = [a_1, a_2, a_3, \dots].$$

Then the periodic approximations corresponding to the sequence of the convergents p_n/q_n to θ' . The relative error E_n of the approximation is no more than $kq_n\|q_n\theta'\|$ (the standard error of $\|q_n\theta'\|$ multiplied by the number of squares, as a single cylinder must wind through all of them, and divided by the width $1/q_n$), and it is standard that

$$\frac{k}{a_{n+1} + 1} \leq kq_n\|q_n\theta'\| \leq \frac{k}{a_{n+1}}.$$

Therefore a good sequence of periodic approximations corresponds to a subsequence of partial quotients diverging, while a fairly good sequence of periodic approximations corresponds to a subsequence of partial quotients at least as large as $2k + 1$, but bounded above. Both situations are clearly satisfied for generic θ' . It is not immediately clear that if M admits a single-cylinder direction that both conditions are generically satisfied along the subsequence of p_n/q_n which correspond to these single-cylinder directions, but this fact holds by the Patterson/Sullivan theorem used in Lemma 2.6.

Theorem 4.3. *Suppose that M is square-tiled and f defines a natural staircase. Assume that d is not a singularity of the flat metric, but is an endpoint of a cut along which f takes the value e_i (and not the end point of any other cut). Assume further that along a subsequence of fairly good periodic approximations there exist*

cylinders $A_{i,j}$ with d on both sides of $A_{i,j}$, and for each of these approximations θ_i there is some P (independent of i) such that for each $x \in A_{i,j}$

$$\|S_{h_{i,j}}(x)\| \leq P,$$

where the sum is with respect to the flow in the periodic direction θ_i . Then e_i is a finite essential value of $\{\tilde{M}, \tilde{\varphi}_t\}$ in direction θ . If d is the endpoint of several cuts, then the essential value is given by the sum of the values of f on the cuts terminating at d . In both cases the staircase is recurrent.




FIGURE 6. A modified orbit-neighborhood used in staircases: $C_{i,1}$ and $C_{i,2}$ are the two colored parallelograms. The region outside the cylinder appears somewhere inside via the appropriate identification and there are no singularities within $C_{i,1}$ or $C_{i,2}$. No cuts have endpoints within these neighborhoods other than the specified cuts with values $\pm e_i$ with common endpoint d . The time h' flow (to traverse the entire vertical distance) acts as horizontal translation by $E_n/q_n \in [\delta/q_n, 1/2q_n]$ on this set. The constant δ is from the definition of fairly good approximation.

Proof. The singularities and ramification points are arranged along the edges of the periodic cylinders $A_{i,j}$. The flow in the direction θ_i induces the rotation by p_i/q_i on the projection of all horizontal bases of the k squares to a single circle. Note that the vertical distance between successive singularities and ramification points on the boundary of the cylinder (i.e. the geodesic distance in direction θ_i between them) must be at least q_i , and the height of any cylinder is no more than kq_i , where k is the number of squares in M . Consider the flow of a vertical interval of length q_i both above and below d on both sides of $A_{i,j}$, and arrange the

cylinder so that this interval is at the bottom of one side. We flow this interval forward until intersecting the height of d within the cylinder, and backwards to the height of the base of the cylinder; this notion is well-defined despite the flow exiting the cylinder in one of these directions as there are no singularities near d (see Figure 6). No other discontinuities of f are in these vertical segments, then, and furthermore the cuts with endpoint d cuts across the cylinder $A_{i,j}$ from exactly one of the two copies of d on the boundary of $A_{i,j}$ (because d is an endpoint of the cut). See Figure 6; the two halves of this orbit neighborhood are $C_{i,1}$ and $C_{i,2}$; because there are no singularities within height q_i of d along the boundary of the cylinder there is no problem extending these neighborhoods outside the boundary. The distances are maintained as drawn, and any cut besides those with endpoint d which might intersect these neighborhoods must pass completely through them; the only endpoint of any cut seen in $C_{i,1}$ and $C_{i,2}$ is d . The difference in the ergodic sums is therefore exactly a sum of $\pm e_i$, the values of f on the cuts terminating at d . As the error E_i is not larger than half the width of $A_{i,j}$ (as we have a fairly good sequence of periodic directions), these neighborhoods may flow for time to traverse the entire height of the cylinder without intersecting the sides (except for the small bit at the bottom where there are no singularities, as mentioned prior). The measure of each orbit neighborhood can be measured by the product of the vertical length of the segment between d and the nearest singularity (which is at least q_i as we are square-tiled) and the horizontal drift; the error term E_i , which is a positive portion of $1/q_i$ as we have assumed only a fairly good sequence of periodic approximations; the measure of each $C_{i,m}$ is at least δ , where $\delta > 0$ is the fixed constant from the definition of a fairly good sequence of periodic approximations.

Therefore the sets $C_{i,1}$ and $C_{i,2}$ again form a quasi-rigidity sequence of sets (their near-invariance under φ_1 is clear from construction). As in Theorem 3.7, the the sum of the values of f on the cuts terminating at d is the difference in sums between infinitely many of the two $C_{i,1}$ and $C_{i,2}$, and is therefore an essential value. \square

Corollary 4.4. *Suppose that for each i , at least one endpoint one of the cuts where f takes value e_i is not a singularity and is not the endpoint of any other cut, and assume that M admits a single-cylinder direction. Then the natural staircase given by f is ergodic in almost every direction.*

Proof. As M admits a single-cylinder periodic direction, almost every θ has a fairly good sequence of single-cylinder periodic directions by Lemma 2.6. As there is only one cylinder to consider, the relevant bound on sums used in Theorem 4.3 is satisfied by Proposition 3.6. As we assume that each cut has a non-singularity endpoint not shared by any other cut, the group of essential values therefore contains each e_i . \square

Example 4.5. Consider the surface on at least five squares given by Figure 7; one row of two squares and one row of at least three. The first presentation with cuts as marked on the left and its associated staircase is *not* ergodic in almost every direction by [8, Thm. 1.4]; The surface is in $\mathcal{H}(2)$ and both cuts are nontrivial in $H_1(M, \mathbb{Z})$. The presentation on the right, however, has both cuts with an endpoint which is not a singularity; by Theorem 4.3, then, the flow on the corresponding staircase is ergodic in almost every direction. Note that the two compact surfaces are the same, and the only difference is in the choice of cuts!

The condition that there is some direction θ' such that $\{M, \varphi_t\}$ decomposes into a single periodic cylinder is convenient but not necessary:




FIGURE 7. Both surfaces are equivalent (given the same number of omitted squares), but the two different realizations as a natural staircase are quite different. Identification is by opposite-side except where marked with a, b .

Example 4.6. Let M be the square-tiled surface in Figure 8, and assume that the number of tiles in each row is odd and at least five. Then if we construct a \mathbb{Z} -staircase using the cuts marked, this staircase is ergodic in almost every direction. This surface, however, admits no single-cylinder directions.

Proof. That each row has more than three squares ensures that the two cuts with value $+1$ do not join together to form a loop; each cut of value $+1$ therefore has both endpoints not singularities. The assumption that the number of squares in each row is odd guarantees the existence of an integer involution point; a vertex around which we may rotate the entire figure by π . Suppose that some flow in slope $q/p > 1$ is a single-cylinder direction. This flow induces a rotation by p/q on the bases of the squares; a periodic cylinder will be of width $1/q$ and may be taken to begin with an interval of length $1/q$ extending horizontally from this involution point. That this surface has no single-cylinder directions follows from Corollary A.5, where this particular example is referenced.

The surface clearly has a decomposition into two cylinders of equal measure and length: the horizontal flow achieves such a decomposition, and the two cylinders are exchanged by the involution. By setting the involution point to be the origin, the involution itself may be represented as $-\text{Id}$. Of course $-\text{Id}$ is in the center of Γ , the Veech group. It follows that any image of these two cylinders under any element of Γ also consists of two components which are exchanged by the involution.

Note that f , the function used to define the staircase, is invariant under this involution. Combining this fact with the fact that the involution exchanges the two cylinders, we see that the ergodic sum in each periodic cylinder is equal to the other. As their sum must be the average of the function, each periodic cylinder sees a total sum of zero. We may then consider the endpoint of the cut with value $+1$ and directly apply the same construction in Figure 6 to find an essential value of ± 1 for this staircase. □

Note that any probabilistic technique of generating covers (not restricted to natural staircases, so the cuts are simple geodesic segments as in §3) on these surfaces which respects the following

- The flow in the horizontal direction of *every* point sees a sum of zero across all cuts traversed through one period,
- the cuts are generated pairwise, invariant under the involution,
- the values taken on the cuts generate the group G




FIGURE 8. The surface M for example 4.6; the labels one through $4n + 2$ on the squares are used in Appendix A. Opposite sides are identified; the surface is in $\mathcal{H}(1, 1)$ and has no single-cylinder directions. We form a staircase using the indicated values on the specified unit intervals.

will generically produce covers which are ergodic in almost-every direction, exactly via the techniques of §3.

APPENDIX A. PERMUTATIONS AND PARALLELOGRAM-TILED SURFACES

While the aforementioned results of Hubert-Lelievre, McMullen, Kontsevich-Zorich and Lanneau-Nguyen guarantee for certain square-tiled M the existence of directions θ_i which decompose M into a single periodic cylinder (providing the bound on ergodic sums necessary in the definition of essential points, Definition 3.5), it is of interest to quickly determine which square-tiled surfaces admit such a direction in general.

Any square-tiled translation surface, with k total squares, may be identified with a pair of permutations in Σ_k (the permutation group of k elements) (σ_h, σ_v) . From a combinatorial perspective it is unimportant that the tiles are squares; any parallelogram tiling would suffice (and still allows a notion of ‘top/bottom’ and ‘left/right’). The permutation σ_h is given by the identification of the right side of the tile indexed by i with the left side of the tile indexed by $\sigma_h(i)$, while the top of tile i is identified with the bottom of the tile of index $\sigma_v(i)$: see Figure 9. The observation that square-tiled surfaces may be indexed by pairs of permutations appears in [27].




FIGURE 9. The permutations σ_h and σ_v used to identify the square-tiled M . σ_h tracks identification of the horizontal flow, and σ_v tracks the vertical flow.

It is clear that any choice of (σ_h, σ_v) along with a choice of parallelogram tile defines a square-tiled translation surface M , and that any choice of charts for a surface M which gives a tiling by parallelograms can be identified with such a pair of permutations. In the search for single-cylinder directions the distinction between squares and parallelograms is unimportant (the two being distinguished only by an affine transformation), so we will use squares without loss of generality.

Theorem A.1. *The parallelogram-tiled translation surface identified by (σ_h, σ_v) is connected if and only if $\langle \sigma_h, \sigma_v \rangle$, the subgroup of Σ_k generated by σ_h and σ_v , acts transitively on \mathbb{Z}_k . The cycle decomposition of the commutator*

$$\tilde{\sigma} = [\sigma_v, \sigma_h] = \sigma_v^{-1} \circ \sigma_h^{-1} \circ \sigma_v \circ \sigma_h$$

encodes the stratum of M : the number of distinct vertices is given by the number of cycles in $\tilde{\sigma}$, and each corresponding vertex has cone angle 2π times the length of the cycle.

Proof. Suppose that there is a proper subset $S \subset \mathbb{Z}_k$ which is invariant under each of σ_h, σ_v . Then the collection of squares indexed by S contains no identification to squares indexed by $\mathbb{Z}_k \setminus S$, thereby forming at least two connected components. On the other hand, if there are no proper invariant subsets, then given any two indices i and j , there is a sequence of σ_{t_n} which connect the square indexed by i to that indexed by j through some sequence of side identifications (where $t_n \in \{h, v\}$).

Consider the upper-right corner of the tile indexed by i . The commutator $\tilde{\sigma}$ exactly encodes the trajectory of a counter-clockwise circle around this corner, and the conclusion is direct from this observation. Let V be the number of cycles in $[\sigma_v, \sigma_h]$. Then M is given as a surface with k faces (the tiles), $2k$ edges (each edge is used in exactly two tiles) and V vertices. By computing the Euler characteristic, the genus g is given by

$$g = \frac{k - V}{2} + 1,$$

assuming that M is connected (otherwise we could consider each connected component separately). \square

Corollary A.2. *Let $\tilde{\sigma}$ be an arbitrary element of the commutator subgroup of Σ_k . Then the number of cycles in $\tilde{\sigma}$ is either odd or even according to whether k is odd or even.*

Proof. It is not difficult to show (for example, by induction on n) that the commutator subgroup of Σ_k is A_k , the alternating subgroup, and that A_k is in fact exactly the set of commutators. The claim may be then shown directly for any element of A_n . However, it may also be derived here from the Euler characteristic and Theorem A.1: the genus must be an integer. In the event that $\tilde{\sigma}$ is the commutator of two permutations with a shared proper invariant subset, the corresponding translation surface is not connected, but we may apply the same reasoning to each connected component. \square

Now suppose that we flow along M in a rational direction (recall that without loss of generality we are using unit squares as tiles). Assume that the slope is $q/p \geq 1$, so that in the time necessary for the flow to traverse a unit of vertical distance, it traverses a horizontal distance of p/q . Then we define a new tiling of M by kq parallelograms: the new tile of index $(i-1)q + j$ (for $i = 1, \dots, k$ and $j = 1, \dots, q$)

is defined to be the parallelogram given by flowing the interval $[(j-1)/q, j/q)$ in the base of the square of index i for a vertical distance of one. See Figure 10. If the flow were of slope less than one, we could perform an analogous retiling using intervals of the vertical edges of the original tiles, flowed to a horizontal distance of one.

The new tiling can also be described with a pair of permutations σ'_h, σ'_v . We are free to choose which direction of these new parallelograms should correspond to ‘vertical’; whether the direction q/p was larger or less than one, there is a natural choice. As our figure uses a flow of slope larger than one, we let the flow correspond to the ‘vertical’ direction in the new tiles.




FIGURE 10. Retiling the surface M based on the flow in direction $q/p \geq 1$. If the large square is the original tile of index i , then the highlighted parallelogram is the new tile of index $(i-1)q + j$.

Lemma A.3. *The new permutations are defined by*

$$\sigma'_h(iq + j) = \begin{cases} (i-1)q + j + 1 & (j < q) \\ (\sigma_h(i) - 1)q + 1 & (j = q), \end{cases}$$

$$\sigma'_v(iq + j) = \begin{cases} (\sigma_v(i) - 1)q + j + p & (j < q - p) \\ ((\sigma_v \circ \sigma_h)(i) - 1)q + j + p - q & (j \geq q - p). \end{cases}$$

Proof. The horizontal identification is direct (refer again to Figure 10). For the vertical identifications, note that flowing to traverse one unit of vertical distance, we either flow only through the top of a square, or first across the right edge and then through the top, exactly according to whether the base of the interval is less than $(1 - p/q)$ or not. \square

These permutations are of limited practical use: the number of tiles increases with the complexity of the slope. We may, however, reduce the problem to a permutation on k elements if we relax the requirements on the parallelograms. In a tiling, it is necessary for all singularities of the metric to belong to the set of

corners of the tiles. Without this requirement, however, we may cover M with k parallelograms with sides horizontal and of slope q/p : the parallelogram of index i is given by flowing the interval $[0, 1/q)$ on the bottom of the square of index i through a vertical distance of q .

It is direct to see that the k parallelograms defined in this way are disjoint: the induced rotation on the bases is of p/q for each unit of vertical distance traversed, so a vertical distance of q is traversed before the left endpoint of this interval intersects an integer point. These tiles must therefore cover M , as we construct k such tiles and each has equal measure to the original squares. It is not a proper tiling, however, as the singularities will presumably belong to the sides of these parallelograms. Now compare with the proper retiling by kq parallelograms, each of the k parallelograms in the improper tiling is tiled by q proper parallelograms, so the proper tiling is a subtiling of the improper tiling. The permutation $\hat{\sigma} \in \Sigma_k$

$$\hat{\sigma}(i) = (\sigma'_v)^q((i-1)q+1)$$

is the permutation induced by this subtiling in the flow direction.

Then the following is clear by our construction:

Theorem A.4. *The flow in slope q/p is a single-cylinder direction if and only if $\hat{\sigma}$ is a single cycle on k elements.*

Generation of $\hat{\sigma}$ from the original pair (σ_h, σ_v) is intimately related to Farey series, the Stern-Brocot tree, and a notion from combinatorics (used mostly in computer science) called *Lyndon words*. On $\{0, 1\}^{\mathbb{N}}$, define the homomorphism ϕ by

$$\phi(0) = \sigma_v, \quad \phi(1) = \sigma_v \circ \sigma_h,$$

and we construct words $\omega_{p/q}$ of length q with exactly p ones in the following manner, beginning with $\omega_0 = 0$ and $\omega_1 = 1$: let $p/q < p'/q'$ and assume that $p'q - pq' = 1$. Let $p''/q'' = (p + p')/(q + q')$ be the *Farey sum* (the ‘bad arithmetic sum’). Then $\omega_{p''/q''} = \omega_{p/q}\omega_{p'/q'}$.

See [20] for background in combinatorics on words; we will not present that background here. A binary word of finite length is called a *Lyndon word* if each left factor has an average strictly smaller than the overall average of the word, i.e. it is strictly smaller in lexicographic order than all of its rotations.

Then our Farey-series-based construction generates all (binary, finite) Sturmian Lyndon words; the correspondence is made quite explicit by considering the geometric construction of Sturmian words by counting crossings of vertical versus horizontal integer-grid segments, and the Lyndon property is quickly seen to be inherited by the given order of composition under the assumption that both $\omega_{p/q}$ and $\omega_{p'/q'}$ are Lyndon words. That of all primitive Sturmian words of a given length and density there is only one which is also a Lyndon word is fairly direct to show; a proof appears in [24].

Note that we defined $\hat{\sigma}$ to track the orbit of the interval $[0, 1/q)$ on the base of a square. Considering a different interval or a vertical segment would amount to performing a fixed cyclic transformation on the sequence of permutations used to define $\hat{\sigma}$. That is, instead of considering

$$\hat{\sigma} = \tau_1 \circ \tau_2 \circ \cdots \circ \tau_n$$

where each τ_i is either σ_h or σ_v , we would consider

$$\tau_k \circ \tau_{k+1} \circ \cdots \circ \tau_n \circ \tau_1 \circ \cdots \circ \tau_{k-1}.$$

This new permutation is naturally seen to be in the same conjugacy class as $\hat{\sigma}$, and therefore has the same number and length of cycles. Similarly, as the labeling of the tiles was arbitrary, it is only the conjugacy class/cycle decomposition of the resulting permutation that is of interest. It would be of considerable interest to classify the permutation pairs (σ_h, σ_v) on k elements which generate a “Farey-Sturmian” composition in the described manner which is a k -cycle. Together with the cycle decomposition of the commutator $[\sigma_v, \sigma_h]$, knowing which pairs can generate k -cycles would classify the square-tiled surfaces with a single-cylinder direction!

Corollary A.5. *If σ_h and σ_v generate a subgroup of Σ_k which contains no k -cycle, then M admits no single-cylinder direction.*

Proof. For slopes larger than one the result is immediate. For slopes less than one, we would perform an analogous construction as indicated previously with the same result. Alternately, any cylinder decomposition for a slope less than one may be mapped via the Veech group to a slope larger than one. \square

Consider for example the surfaces given in Figure 8, with squares labeled as in that figure:

$$\begin{aligned}\sigma_h &= (1 \dots (2n+1))(2n+2) \dots (4n+2), \\ \sigma_v &= (1)(2(2n+2))(3(2n+3)) \dots ((2n+1)(4n+1))(4n+2).\end{aligned}$$

Both σ_h and σ_v are in the alternating group A_6 (σ_h is the product of two disjoint cycles of equal length, and σ_v is explicitly given as an even number of transpositions), but a $4n+2$ -cycle is a composition of $4n+1$ transpositions. Therefore by the above corollary, *this surface admits no single-cylinder directions.*

Example A.6. The *eierlegende Wollmilchsau* is the surface on eight squares defined by

$$\sigma_h = (1234)(5678), \quad \sigma_v = (1836)(2745).$$

This surface is of genus 3 with four singularities of cone angle 4π : from the permutation presentation, the number of vertices (and therefore the genus) and their cone angles may be read off the commutator of σ_v and σ_h :

$$[\sigma_v, \sigma_h] = (13)(24)(57)(68).$$

Both σ_h and σ_v are in A_8 , while a complete 8-cycle has seven transpositions. Therefore this surface has no single-cylinder directions. This result also follows from the fact that the Veech group is all of $\mathrm{SL}_2(\mathbb{Z})$ [10]. Therefore all periodic decompositions are of the same type: two cylinders of equal heights and measures.

Corollary A.7. *Let σ_v and σ_h define a (connected) square-tiled surface, and let $\tilde{\sigma} = [\sigma_v, \sigma_h]$. If $\tilde{\sigma}$ contains at least d cycles of length one, then it is possible to construct a \mathbb{Z}^d natural staircase whose geodesic flow is ergodic in almost every direction. On the other hand, if there are no cycles of length one and $M \in \mathcal{H}(2)$, then no natural staircase is ergodic in almost every direction.*

Proof. The statement about the possibility of ergodic d -dimensional natural staircases follows from Corollary 4.4, using the vertices which are not singularities (fixed points under $\tilde{\sigma}$) as endpoints of distinct cuts whose values generate \mathbb{Z}^d . If there are no cycles of length one, then all integer points are singularities. All cuts are therefore nontrivial in the relative homology group $H_1(M, D, \mathbb{Z})$. In the event that $M \in \mathcal{H}(2)$, then, the geodesic flow is not ergodic in almost every direction, again citing [8]. \square

The above techniques may be used to quickly generate surfaces with single-cylinder directions. Assuming that a surface is tiled by parallelograms one through k such that vertical flow is a single cylinder, the permutation σ_h therefore encodes the stratum of the surface. Letting the tiles be given in order, so that $\sigma_v(i) = i + 1 \bmod k$, the singularities may be deduced from the cycle decomposition of the commutator

$$[\sigma_v, \sigma_h](i) = \sigma_h^{-1}(\sigma_h(i) + 1) - 1.$$

For arbitrary dimension d the techniques of this appendix may be used to study d -cube tiled translation manifolds through the use of d permutations (for the relevant face identifications in each dimension). A single-cylinder direction therefore gives rise to a parallelepiped tiling for which one of the permutations is a single cycle. The geometric data is still encoded by the permutations, if more involved to extract. For example, the number of connected components still corresponds directly to the number of invariant sets in the group generated by the given permutations, and the singularities of the flat metric at certain $(d - 2)$ -dimensional simplices involves commutators of the permutations. As geodesic flows in higher dimensions are poorly understood at present, we have phrased all the above in the particular case $d = 2$.

REFERENCES

- [1] J. Aaronson. *An introduction to infinite ergodic theory*, volume 50 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1997.
- [2] M. Boshernitzan. A criterion for weak mixing of induced interval exchange transformations. *ArXiv e-prints*, may 2011.
- [3] J. Chaika and P. Hubert. in preparation. preprint.
- [4] J.-P. Conze. Recurrence, ergodicity and invariant measures for cocycles over a rotation. *Contemporary Mathematics*, 485:45–70, 2009.
- [5] J.-P. Conze and K. Frączek. Cocycles over interval exchange transformations and multivalued Hamiltonian flows. *ArXiv e-prints*, mar 2010.
- [6] J.-P. Conze and E. Gutkin. On recurrence and ergodicity for geodesic flows on noncompact periodic polygonal surfaces. preprint.
- [7] V. Delecroix, P. Hubert, and S. Lelièvre. Diffusion for the periodic wind-tree model. preprint.
- [8] K. Frączek and C. Ulcigrai. Non-ergodic \mathbb{Z} -periodic billiards and infinite translation surfaces. *ArXiv e-prints*, sep 2011.
- [9] E. Gutkin and C. Judge. Affine mappings of translation surfaces: geometry and arithmetic. *Duke Math. J.*, 103(2):191–213, 2000.
- [10] F. Herrlich and G. Schmithsen. An extraordinary origami curve. *Mathematische Nachrichten*, 281(2):219–237, 2008.
- [11] P. Hooper and B. Weiss. Generalized staircases: recurrence and symmetry. preprint.
- [12] W.P. Hooper, P. Hubert, and B. Weiss. Dynamics on the infinite staircase. *Disc. Cont. Dyn. Sys.* to appear.
- [13] P. Hubert and S. Lelièvre. Prime arithmetic Teichmüller discs in $\mathcal{H}(2)$. *Isr. J. Math.*, 151:281–321, 2006.
- [14] P. Hubert, S. Lelièvre, and S. Troubetzkoy. The ehrenfest wind-tree model: periodic directions, recurrence, diffusion. *Journal fuer die reine und angewandte Mathematik (Crelle's Journal)*, 656:223–244, 2011.
- [15] P. Hubert and T. A. Schmidt. An introduction to Veech surfaces. In *Handbook of dynamical systems. Vol. 1B*, pages 501–526. Elsevier B. V., Amsterdam, 2006.
- [16] P. Hubert and B. Weiss. Ergodicity for infinite periodic translation surfaces. preprint.
- [17] S. Kerckhoff, H. Masur, and J. Smillie. Ergodicity of billiard flows and quadratic differentials. *The Annals of Mathematics*, 124(2):pp. 293–311, 1986.
- [18] M. Kontsevich and A. Zorich. Connected components of the moduli spaces of Abelian differentials with prescribed singularities. *Invent. Math.*, 153(3):631–678, 2003.
- [19] E. Lanneau and D.-M. Nguyen. Teichmüller curves generated by weierstrass prym eigenforms in genus three and genus four. preprint.

- [20] M. Lothaire. *Combinatorics on words / M. Lothaire ; foreword by Roger Lyndon*. Addison-Wesley, Advanced Book Program, World Science Division, Reading, Mass. :, 1983.
- [21] H. Masur. Ergodic theory of translation surfaces. In *Handbook of dynamical systems. Vol. 1B*, pages 527–547. Elsevier B. V., Amsterdam, 2006.
- [22] Curtis T. McMullen. Teichmüller curves in genus two: Discriminant and spin. *Mathematische Annalen*, 333:87–130, 2005. 10.1007/s00208-005-0666-y.
- [23] S. J. Patterson. Diophantine approximation in Fuchsian groups. *Philos. Trans. Roy. Soc. London Ser. A*, 282(1309):527–563, 1976.
- [24] D. Ralston. Heaviness in symbolic dynamics: substitution and Sturmian systems. *DCDS-S*, 2:287–300, 2009.
- [25] K. Schmidt. *Cocycles of ergodic transformation groups*, volume 1 of *Lecture notes in mathematics*. MacMillan, India, 1977.
- [26] K. Schmidt. A cylinder flow arising from irregularity of distribution. *Compositio Mathematica*, 36(3):225–232, 1978.
- [27] G. Schmithüsen. Origamis with non congruence Veech groups. In *Proceedings of 34th Symposium on Transformation Groups*, pages 31–55. Wing Co., Wakayama, 2007.
- [28] D. Sullivan. Disjoint spheres, approximation by imaginary quadratic numbers, and the logarithm law for geodesics. *Acta Math.*, 149(3-4):215–237, 1982.
- [29] S. Troubetzkoy. Recurrence in generic staircases. 23:1047–1053, 2011.
- [30] W. Veech. Boshernitzan’s criterion for unique ergodicity of an interval exchange transformation. *Erg. Thry. Dyn. Sys.*, 7:149–153, 1987.
- [31] A. Zorich. Flat surfaces. In *Frontiers in number theory, physics, and geometry. I*, pages 437–583. Springer, Berlin, 2006.

BEN GURION UNIVERSITY, DEPARTMENT OF MATHEMATICS, POB 653, BEER SHEVA, 84105, ISRAEL

E-mail address: `ralston.david.s@gmail.com`

AIX-MARSEILLE UNIVERSITY, CPT, IML, FRUMAM, MARSEILLE

E-mail address: `troubetz@iml.univ-mrs.fr`